VA-118-001-06-TR

Adaptive and Iterative Processing Techniques for Overlapping Signatures

Technical Summary Report

March 2006

PERFORMING ORGANIZATION

AETC Incorporated
1225 South Clark Street, Suite 800
Arlington, VA 22202

This research was supported wholly by the U.S. Department of Defense, through the Environmental Security Technology Certification Program (ESTCP) through Project 200415 under Contract W912HQ-04-C-0017.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
MAR 2006	**Final**	**-**

4. TITLE AND SUBTITLE
Adaptive and Iterative Processing Techniques for Overlapping Signatures

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Tom Bell

5d. PROJECT NUMBER
MM-0415

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
SAIC (formerly AETC Inc.) 1225 South Clark Street, Suite 800 Arlington, VA 22202

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Environmental Security Technology Certification Program 901 North Stuart Street, #303 Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
ESTCP

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**UU**	**19**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

Table of Contents

1. Background and Goals

A primary goal of the UXO research community is to develop technologies that detect and localize buried UXO, and that discriminate them from clutter. Electromagnetic sensors such as EMI systems operate by detecting the presence of an anomalous electromagnetic field that could be caused by buried UXO. Physics-based modeling and analysis procedures, developed under previous SERDP and ESTCP funding for electromagnetic induction and magnetometer sensor data of isolated targets, have been shown to discriminate UXO from clutter based on the derived source parameters for spatially discrete target signatures. However, many real world UXO remediation sites contain highly-contaminated regions with high density of anomalies, both UXO and clutter. In these cases, where the signatures from multiple targets overlap, whether or not they are UXO, the standard procedures do not work well.

The primary problem with overlapping signatures is that conventional inversion procedures assume a single source whose signature is spatially separated from other signatures. Currently, analysts attempt to isolate anomalies by carefully selecting the data to be inverted (usually by manually inspecting a two-dimensional plot of the anomaly and carving out two separate regions of data) and assuming that the selected data reflect the signature caused by a single source. This is time-consuming, requires a good deal of experience, and is not very effective. Better methods are needed to extract the individual target parameters from a multiple target signature.

The goal of this project was to develop advanced iterative techniques for inverting magnetic and electromagnetic data for situations in which the signatures from two targets overlap. After developing the methodology, the algorithm(s) would be tested first on synthetic data without any added noise. The synthetic data would be used to systematically vary the parameters of the two targets by changing their depths, the distance between them and their relative orientations. Later, controlled test data would be used to further validate and test the algorithms in real-life situations.

3. Technical Approach

3.1 General Inversion Methodology

For UXO detection and discrimination, models are fit to survey electromagnetic or magnetometer data using indirect (iterative) inverse methods. These all follow the common process in which (1) model parameters are input to a forward model, (2) output is compared against observed data and used to find a figure of merit (commonly chi-squared error), which is (3) used to generate a new set of model parameters (guesses) for the next iteration. Successive iterations lead to improved fits, until an optimum fit (locally optimum, at least) is found. Because this iterative inversion is performed many times, a fast algorithm is desired.

electromagnetic data Our method for electromagnetic data, like almost all inversion routines, uses the dipole approximation. Each target is completely characterized by three positions (x_o,y_o,z_o), three angles (φ,θ,ψ), and 3*N_t betas ($\beta_1(t_i)$, $\beta_2(t_i)$, $\beta_3(t_i)$), where N_t is the number of time gates for the sensor. The betas are the eigenvalues of the symmetric effective magnetic polarizability tensor, and represent the response of the target along its three principal axes.

We can reduce the number of fit parameter by making use of the fact that the modeled sensor response is linear in the betas. We perform a non-linear Levenberg-Marquardt [1] inversion on the position and angles, with an embedded linear determination of the β's at each iteration. We note that this method produces one set of positions and orientations for all time gates. The algorithm continues until the χ^2 fit between the predicted and measured response at successive iterations changes by less than a set tolerance.

Initial guesses must be provided for the 6 spatial parameters. We set all three angles to 45°. The position parameters x_o and y_o are determined from a signal-weighted mean of the measurement locations. Previous experience with this method has shown that the final results depend most strongly on z_o, since some initial choices for z_o lead to local, rather than global, minima. It was therefore decided to loop over several different initial guesses for z_o that covered a reasonable range for the targets. The results with the smallest mismatch between measured and modeled data (i.e., best χ^2) is chosen as the solution.

magnetometer data Two methods were considered for the inversion of magnetic data. In both cases, the target is characterized by three positions (x_o,y_o,z_o), two angles (dec,inc) for the orientation of the induced dipole (assuming an axisymmetric target), and the radius (a) of an equivalent sphere. In the first method, the algorithm first fits the shape of the footprint, then its magnitude. Initial guesses for the fit parameters are determined internally within the code based on the measured signature. Tests on synthetic, noise-free data did not produce excellent results for this method. This second method is a *two-stage approach*, in which output from the first stage is used as input for the second stage. The first stage provides a coarse estimate of target location using a simplification of the problem, which cuts the number of fitted variables in half, and the second stage provides a more precise solution using the full parameterization of the problem. The presence of the first stage effectively makes the algorithm robust against local minima and poor guesses for the initial parameter values, which, unlike the first algorithm, are provided by the user rather than calculated internally.

The first stage uses the assumption that the dipole field D from the target is much weaker than the Earth's field $\mathbf{B_E}$ at all measurement locations (valid only if the sensor is not too close to the target), so that the magnitude of their vector sum is approximately equal to $|\mathbf{B_E}|$ plus the component of D that lies in the direction $\mathbf{B_E}$. This allows a linear solution of dipole strengths along the coordinate axes so that the fitted parameters are reduced to only the 3 positions. Within each iteration of the loop, the x-, y-, and z-components of the best-fit target dipole (M_x, M_y, M_z) are determined through linear regression on the data. At the end of the process, the best-fit inclination and declination are derived from (M_x, M_y, M_z), and fed into the 2nd stage as a starting point.

The second stage executes a non-linear search on six parameters: the three positions, two angles, and a scalar offset applied to the whole data set to account for various problems (e.g., drift in the earth's field strength or local geologic anomalies). The magnitude, and therefore, the equivalent sphere radius, is determined through scaling. This two-stage method was found to give better inversion results on noise-free synthetic data than the first method, and was used to produce all the magnetic results discussed for this project.

3.2 Two-Dipole Fit Methods

The general inversion methodology described above is followed for any data, whether fitting a single target or multiple targets. For data that potentially contains two dipole signatures, however, the fit algorithm must account for the multiple sources. We developed a two-dipole iterative residual routine for this fit and then, to improve on the performance, developed a simultaneous two-dipole fit algorithm (called "double happiness").

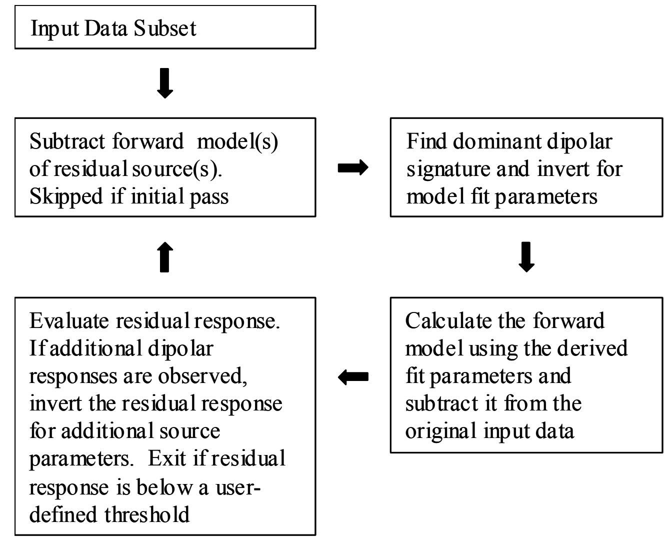

Figure 1. Logic flow for the two-dipole iterative residual fit algorithm.

3.2.1 Two-Dipole Iterative Residual Fit Procedure

We assume that two dipoles produce an overlapping signature that can be generated by adding individual anomalies (i.e., there is no coupling between the targets). The *iterative residual* fit routine is a straightforward approach to disentangling an overlapping signature by decomposing the measured response field into two dipole sources by multiple sub-iterative subtractions, as shown in Figure 1. The anomaly data is inverted to determine the dominant dipolar source and model parameters that represent this dominant source are estimated. A forward model using these estimated parameters is subtracted from the original data and the residual is examined for the presence of a second source. If one is detected above the background noise, the residual is fit as the second dipole, its modeled signal subtracted from the total anomaly, and this new residual is again fit as the first dipole. This basic loop is repeated until the summed fit of the two dipoles matches the total anomaly to within some specified tolerance.

As will be seen in Section 4, this method was only moderately successful on synthetic data without noise, for objects with lateral separations of at least 40 cm, and it suffered from two shortcomings. First, the nature of the method tends to favor cases with one strong and one weak dipole; two nearly equal dipoles are much harder to fit accurately. Second, the method is relatively slow, since there are multiple nested iterations required. A different method, described in the next section, was investigated that proved to be more effective.

3.2.2 Double Happiness Algorithm

This approach developed a simultaneous two-dipole inversion (called "double happiness"), which does not require the extra iterations of the two-dipole iterative residual method and is therefore much faster. Here, the measured anomaly is fit directly by two dipoles, with the fit iterating to convergence. The double happiness algorithm is applied in this project to both electromagnetic and magnetometer data.

The approach taken by this algorithm for electromagnetic data is to express the target response as a tensor, representing the coefficient of proportionality between the illuminating field and the resulting induced dipole. Since the predicted signals are linear with elements of the response tensor, we can linearize most parameters and solve for them directly in each iteration of the search loop. Nonlinear

search is applied only on target locations (x, y, z for both targets; 6 parameters total), and the elements of both response tensors are solved for directly in each iteration. This assumes only the dipole model approximation. The tensors are symmetric, so six values define each completely and the target orientation angles and beta values (response coefficients) are expressed within the elements of this tensor. For a two-dipole solution, 12 parameters are found using linear methods in each iteration. After iteration is complete, rotation matrices are found that diagonalize the tensors, revealing target orientations. This approach cuts parameters in the nonlinear search down to 6. A similar approach is used for magnetometer data.

An interesting side effect of this approach is that the rotation matrices may define different dipole orientations for different time gates. This can be a physically real effect for an isolated target, since different parts of some targets at some orientations can have different time decays, but can also be an indicator of overlapping targets within the signature.

4. Performance on Synthetic Data

4.1 Synthetic Data Generation

Synthetic data sets for testing the two-dipole fit algorithms were created using a Monte Carlo approach. The advantage of a Monte Carlo method is that any number of target types, target combinations, target orientations, relative distances and depths can be chosen for analysis and comparison. We randomly chose two targets from a pool of six UXO targets and varied their separation, depths and orientations.

All data sets were generated on a 9m x 9m grid, with a transect spacing of 0.25m in both x and y. The targets were created using the forward dipole models of the inversion algorithms. The models were based on a total field magnetometer and an EM61 sensor. (The electromagnetic runs used an EM61MkI for the iterative residual method and an EM61MkII for the double happiness algorithm. This can affect the performance comparisons, as discussed later.) The size of the magnetic targets, and the betas for the EM targets, were chosen to roughly correspond to 20mm, 40mm, 57mm, 81mm, 105mm, and 155mm ordnance. For each data set, two of the target types were randomly paired. The target depths were random, but were constrained to fall between 0 and the rule of thumb for the maximum penetration depth for that target (11 times the diameter) as determined by the CORPS DID OE-005-05.01. The target orientations were random and unconstrained. Target separations in the xy-plane were distributed from 0.1m to 2m, equally divided among 4 bins (0.1-0.3m, 0.3-0.7m, 0.7-1.2m, and 1.2-2m). No noise was added. For electromagnetic and magnetic cases, 2000 data sets were produced, with 500 in each lateral separation bin.

4.2 Performance of the Iterative Residual Method

We applied the iterative residual method to the 2000 synthetic data sets described in the previous section. Preliminary testing showed that while one could use the entire grid for inversion of magnetic targets, it was often necessary to manually specify subsets of the grid for the more widely separated EM targets (i.e., draw polygons in a manner similar to the traditional method). This is because the algorithm has difficulty when it is forced to fit two roughly equal targets with only one dipole (in the first iteration). This may be a serious problem with the method and, if it performs well, must be addressed in the future. Because of this, only 80 cases were used for the EM data.

Figures 2 and 3 show the synthetic data results for the iterative residual method on magnetic and electromagnetic data, respectively. Shown are plots of the errors (averaged over both dipoles) in the derived target parameters (x, y, depth, relative size and enclosed solid angle), as a function of the 3D separation of the targets. The total coherence for the two-dipole fit to the data is also shown. The size parameter shown for the EM fits is derived empirically for the EM61MkI and is given by Size=$0.05*(\beta_1+\beta_2+\beta_3)^{1/3}$.

Figures 4 and 5 plot the same errors versus the lateral separation instead of the 3D separation. Differences are slight, but they show that the errors for 3D separation greater than 1m are due to depth differences rather than horizontal separation, which agrees with our physical understanding of the problem. The iterative residual method starts to have difficulty at separations about 1m, although there are many cases that give good results at much smaller separations.

We also tested the iterative residual algorithm on an artificial signal created by taking magnetic and electromagnetic measurements in air of an isolated horizontal 81mm target and an isolated vertical, nose-up 155mm target for separation, adding them together to create artificial magnetic and electromagnetic overlapping signatures, and then increasing the lateral separation in increments of 0.25m. Figure 6 demonstrates the performance of the iterative residual method on magnetic and electromagnetic artificial signals as a function of lateral separation of the two targets, and also compares it to the traditional method, in which polygons are defined manually to delimit the range of each target's footprint. The plots show the fitted positions and the total coherence of the two-dipole fit to the data as a function of lateral separation. For the magnetic signal, the fitted size is proportional to the amplitude of the signal. For the EM plots, the beta sum is formed by simply adding the three betas, and is an indicator of target size. The beta aspect ratio is formed by dividing the longitudinal beta by the transverse beta, and is an indicator of target shape. In each plot, the correct parameter values are plotted as dashed lines, and the derived values as squares. In the results for the traditional method, green squares are used when the targets are too close to discern the presence of two objects and the fit is for one dipole.

For this pair of targets, the iterative residual algorithm performs much better than the traditional method. The latter cannot distinguish overlapping dipoles at lateral separations less than one meter, and even for separations larger than one meter, especially for the EM signal, the derived values are not very accurate. The iterative residual method, on the other hand, can distinguish the two dipoles at lateral separations less than 50cm. However, as expected, the iterative residual inversion performs poorly for very small separations, where the two signatures begin to merge into one and it becomes difficult for the algorithm to distinguish the two dipoles. Because the iterative residual method was tested on EM61MkI synthetic data, which has measurements at a single time gate, it was not unexpected that it had trouble distinguishing the two dipoles at very close ranges.

4.3 Performance of the Double Happiness Method

The double happiness algorithm was developed as an alternative method for discrimination of multiple targets at close separations. Results of this algorithm are shown in Figures 7 and 8 for magnetic and electromagnetic synthetic data, respectively. The same errors in x, y and depth are plotted, as well as the overall coherence of data with the two-dipole solution. For the EM data, response betas were fixed for

each target and the percent of correct solutions are plotted as a function of 3D separation; they range from 86% to 100%.

Comparing Figures 2 and 7, and Figures 3 and 8, it is clear that, for synthetic data without any added noise, the double happiness method performs excellently, and much better than the iterative residual method. Ground truth target parameters were derived with exact agreement in 98% of the cases), although there is still a trend of worsening results with decreasing 3D separation.

The difficulty with the comparison of results for the iterative residual and double happiness methods on EM synthetic data (there is no problem for magnetometer data) is that the former was tested on an EM61MkI data set and the latter on an EM61MkII data set. Since the MkII version of the EM61 takes measurements at three time gates, it is possible that this improves the performance of the algorithm. However, this is a potential, partial, explanation of the results only if the later time gates contain sufficient signal to affect the fits. Our experience has been that the signal falls off fast enough with time that this is not often the case, unless the time gate response is weighted toward the later measurements.

While performance results as a function of horizontal separation are not available, we expect it to improve, especially the EM case, equivalent to the change in the iterative residual algorithm.

Results of the Two-Dipole Iterative Fit algorithm using the new 2-Stage fitter.
* Same set of 2000 cases.
* Average solve time: 18 s.

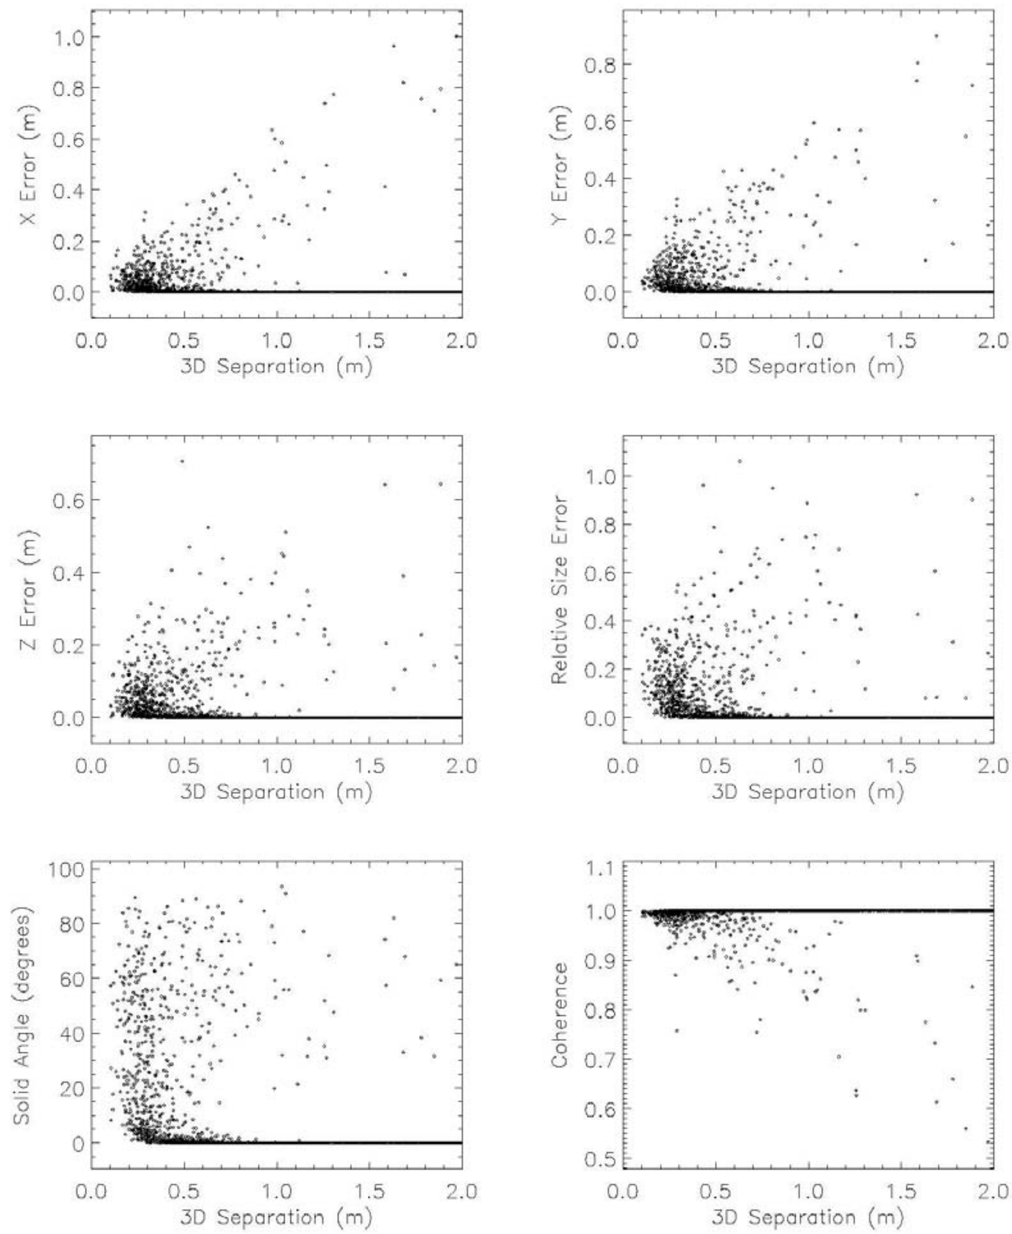

Figure 2. Two-Dipole Iterative Residual method on *magnetic* synthetic data, 2000 cases. All quantities are plotted versus *3D* separation.

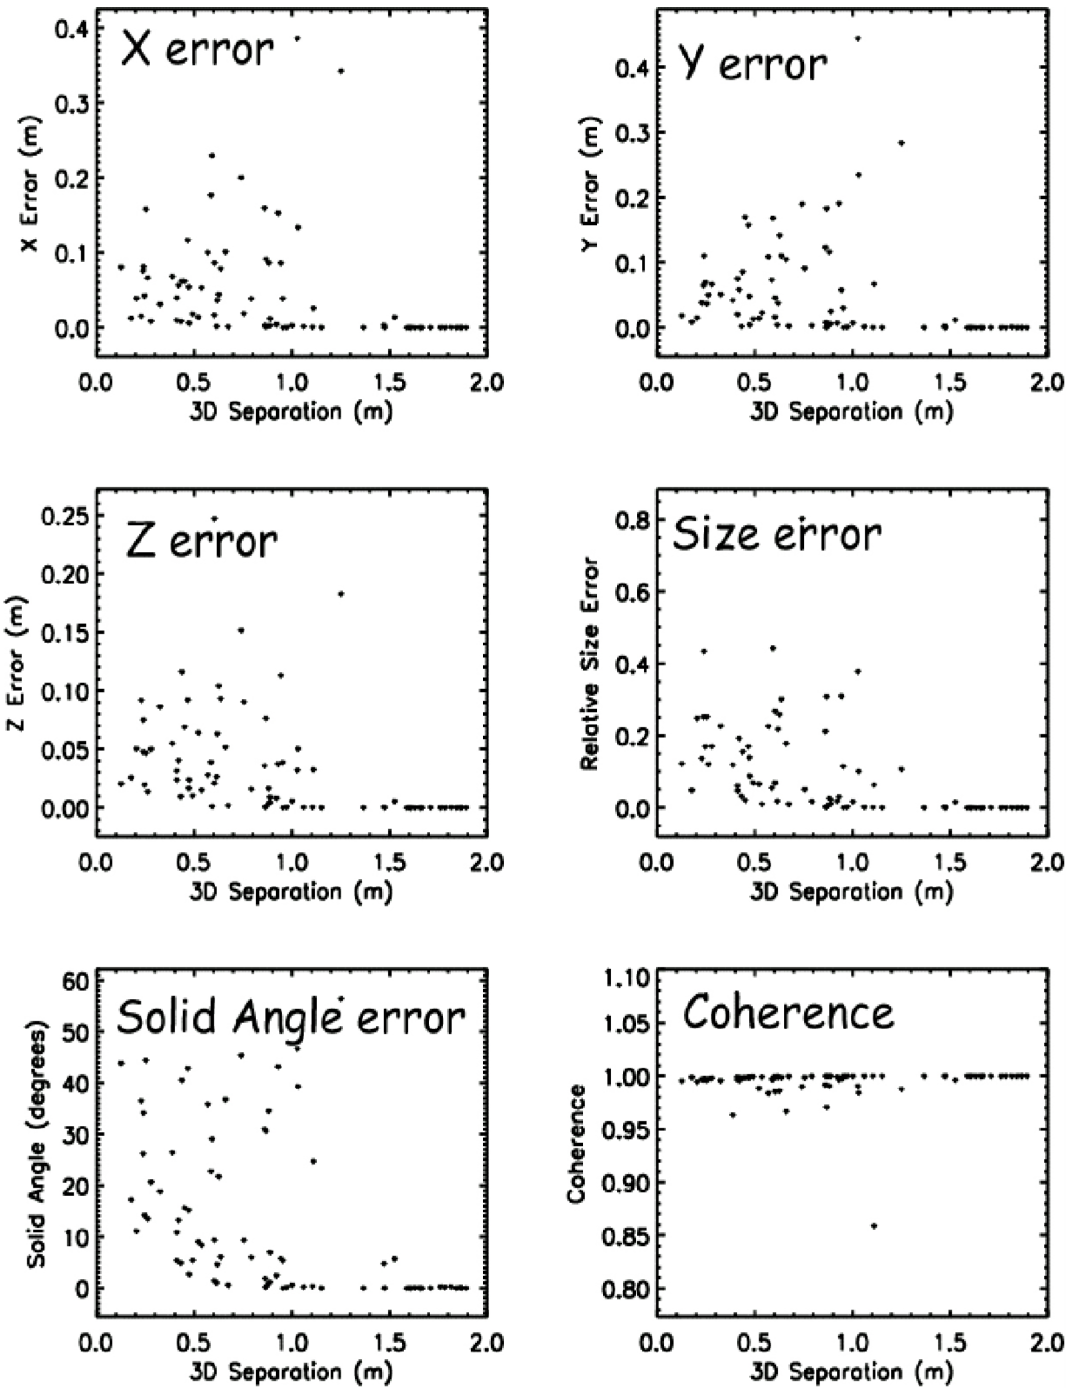

Figure 3. Two-Dipole Iterative Residual method on *electromagnetic* data, 80 cases. All quantities are plotted versus *3D* separation.

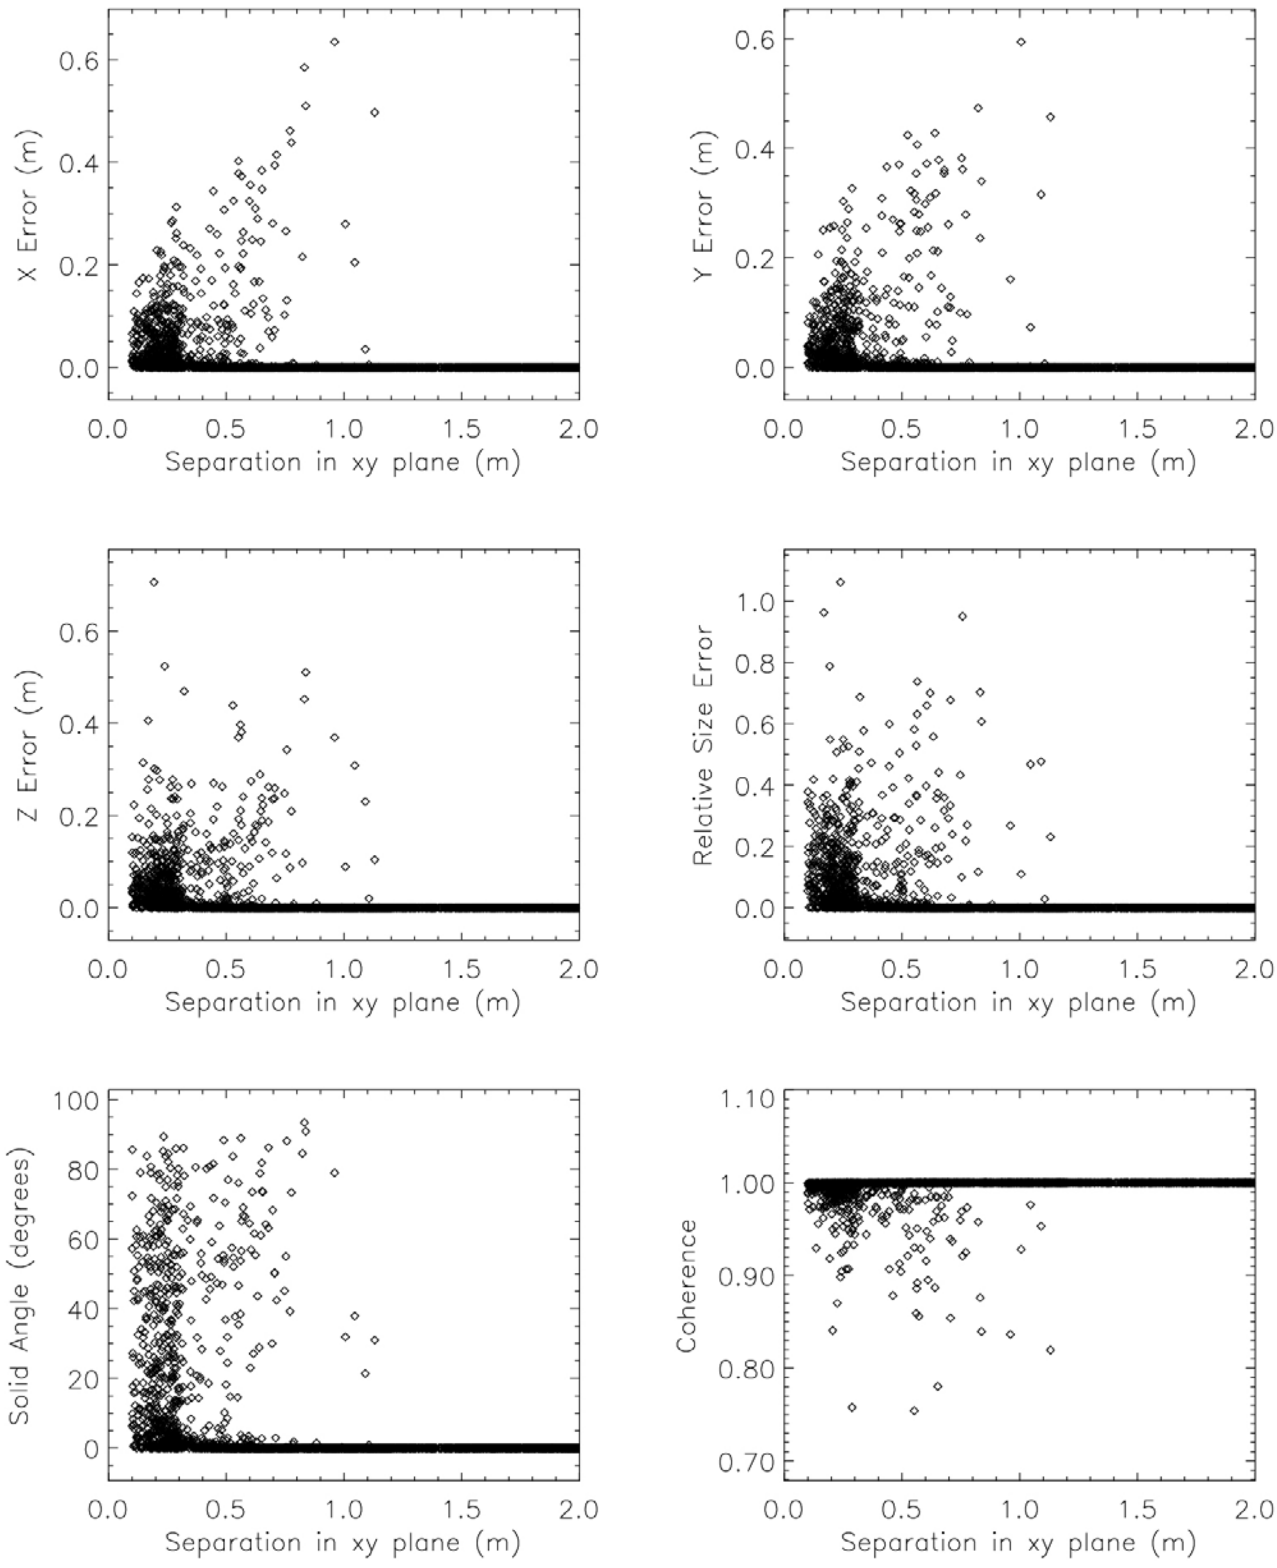

Figure 4. Two-Dipole Iterative Residual method on *magnetic* synthetic data, 2000 cases. All quantities are plotted versus separation in the *horizontal* plane.

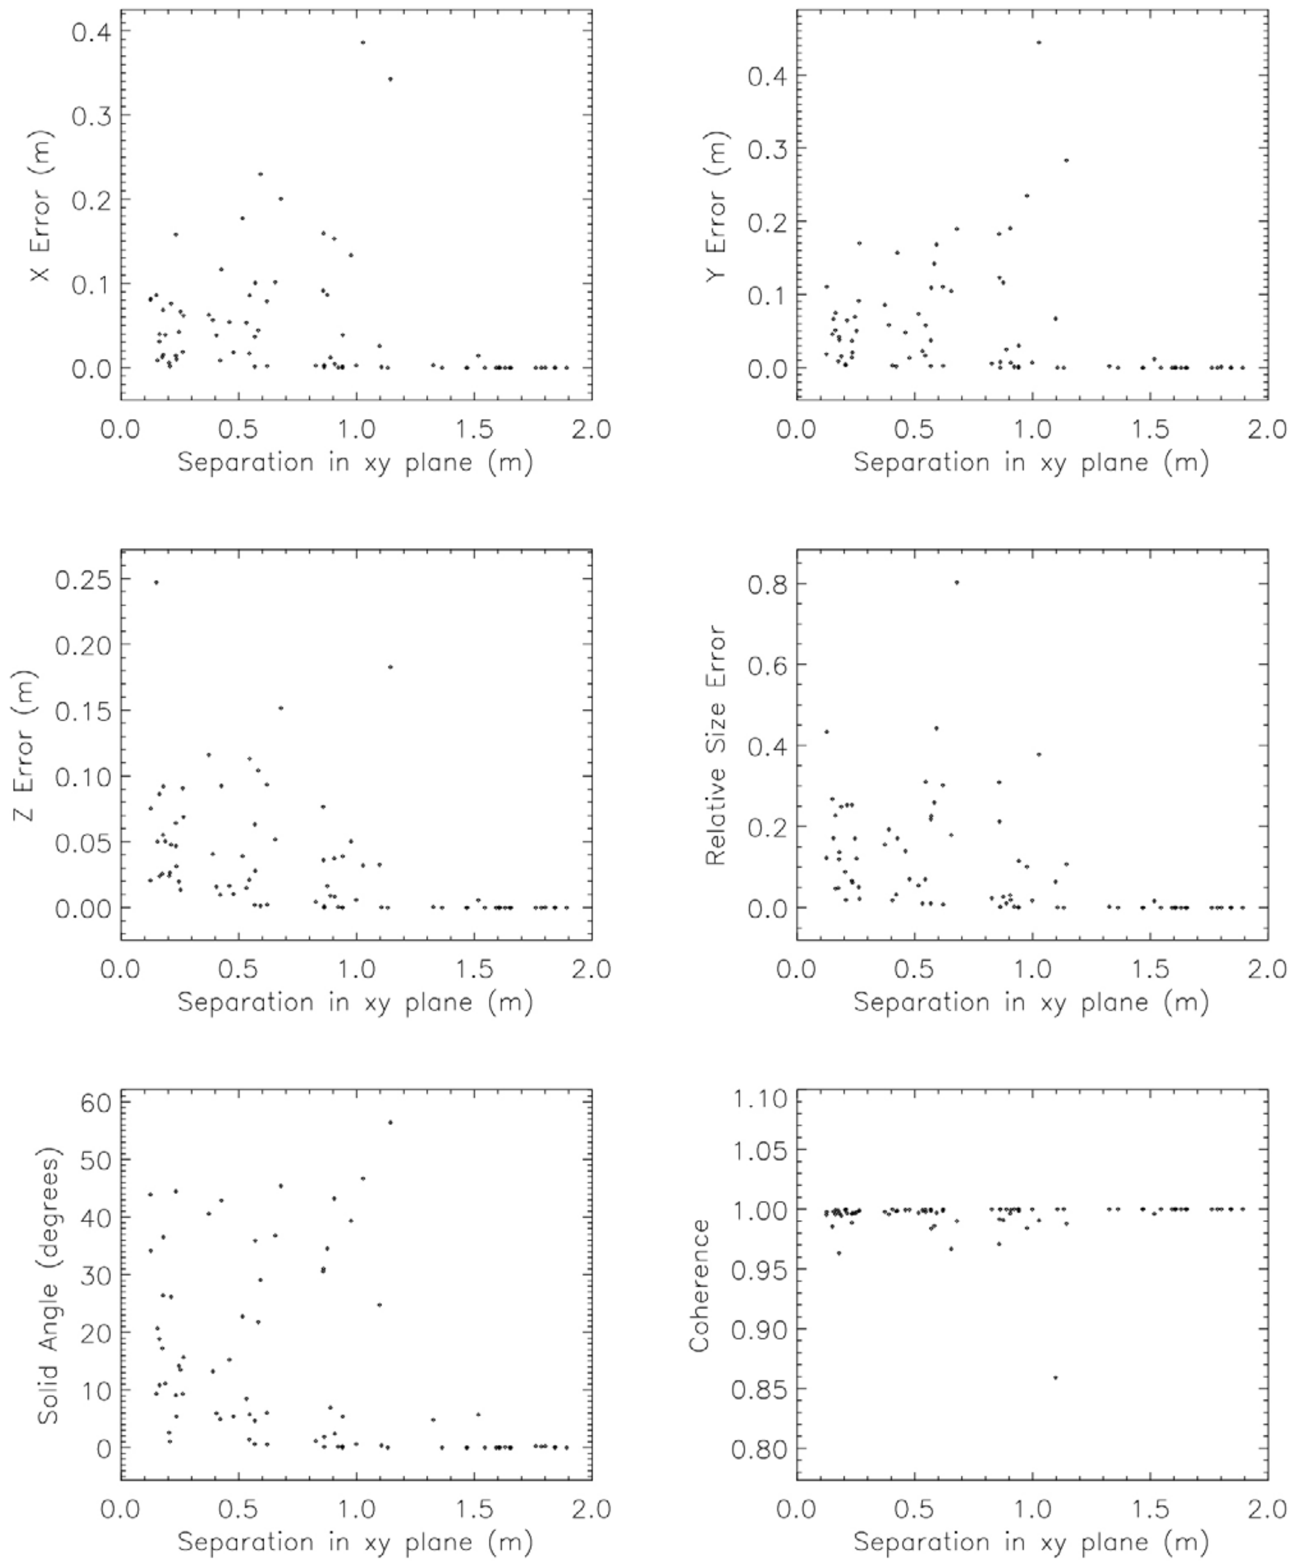

Figure 5. Two-Dipole Iterative Residual method on *electromagnetic* synthetic data, 80 cases. All quantities are plotted versus separation in the *horizontal* plane.

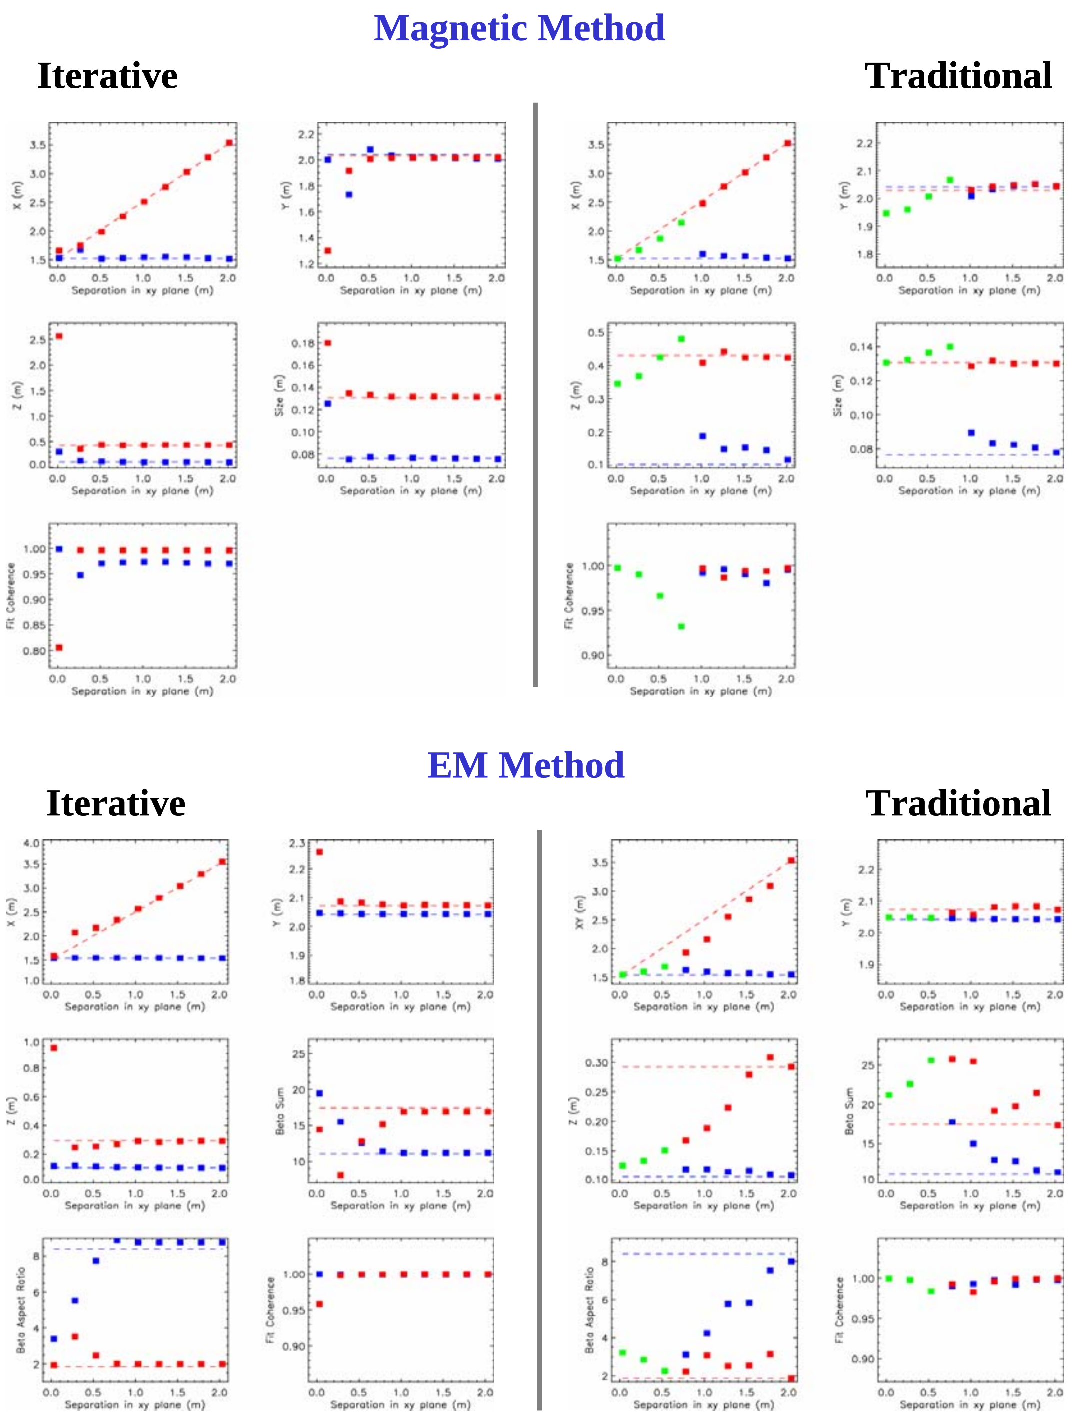

Figure 6. (left) The performance of the iterative residual method on a single artificial signal from two UXO as a function of horizontal separation. (right) The performance of the traditional method of manual definition of polygons to delimit the range of each target.

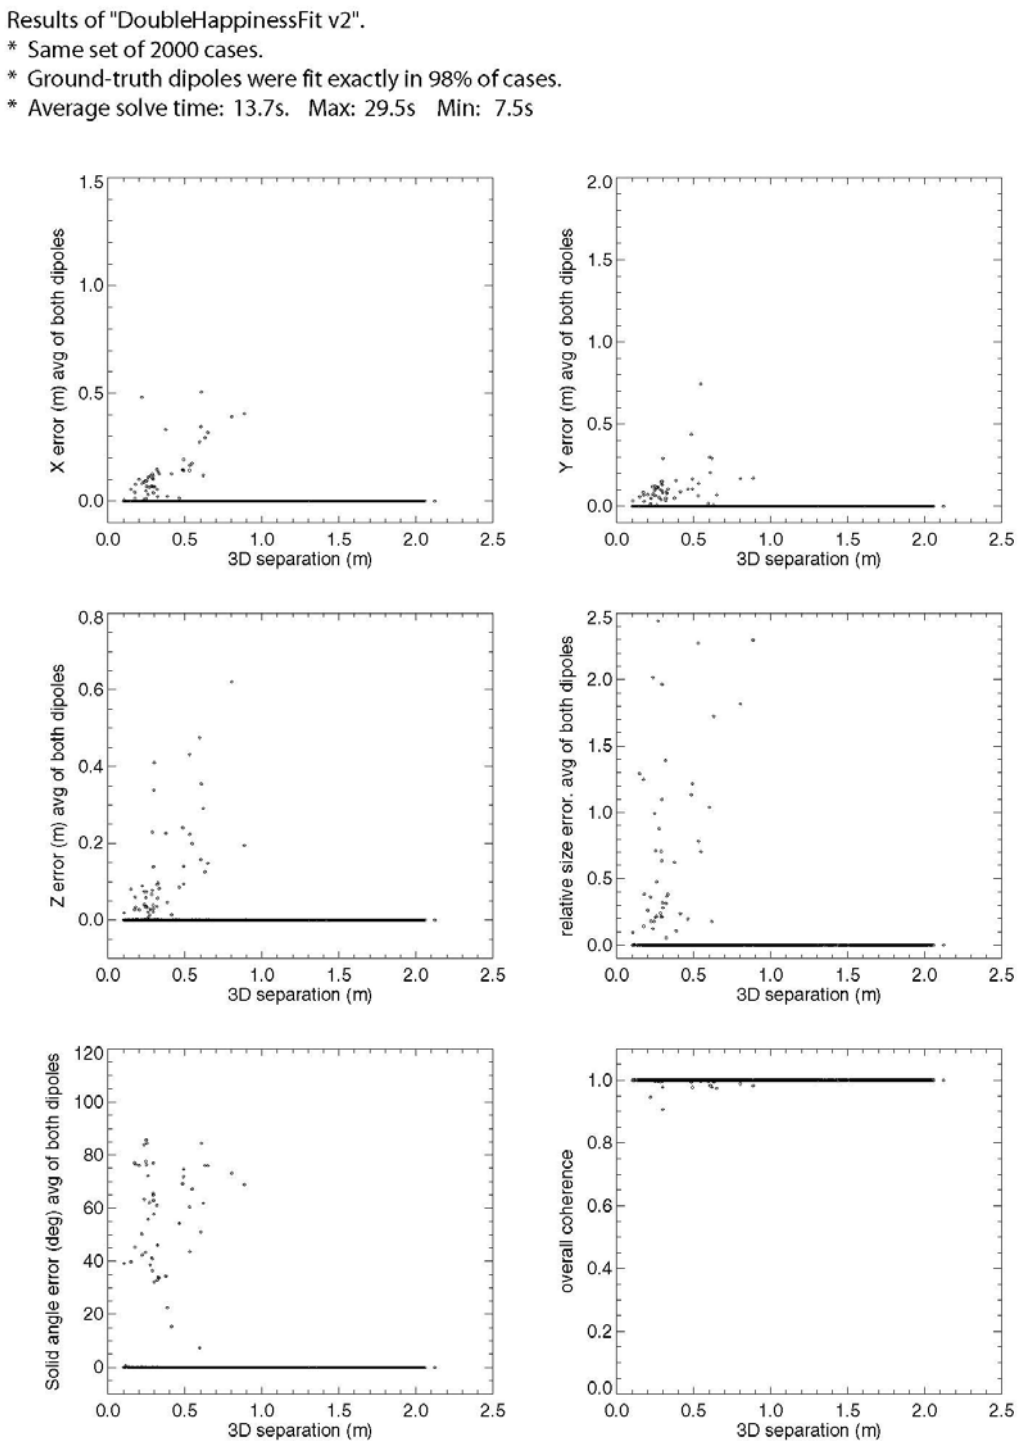

Figure 7. Double Happiness method on *magnetic* synthetic data, 2000 cases. All quantities are plotted versus separation in the *3D* plane.

Results of fitDoubleEM() on synthetic EM61mkII data. Two targets, overlapping signatures.

* 2000 random cases. Six UXO targets formed the pool from which random samples were drawn: {20mm,40mm,57mm,81mm,105mm,155mm}. Response betas fixed for each target.
* Data projected onto 9m x 9m grid with 25cm separation. NO NOISE.
* XY positions random inside 4m by 4m central region. Z ranges from 0 to 11 * diameter.
* Target separations in 3D uniformly distributed from 0.1 to 2.1 m.
* Overall, ground-truth correctly recovered in 93.9% of cases.
* Solve times: 4.6s (min) 45.8s (avg) 165.3s (max) on Pentium 4 desktop.

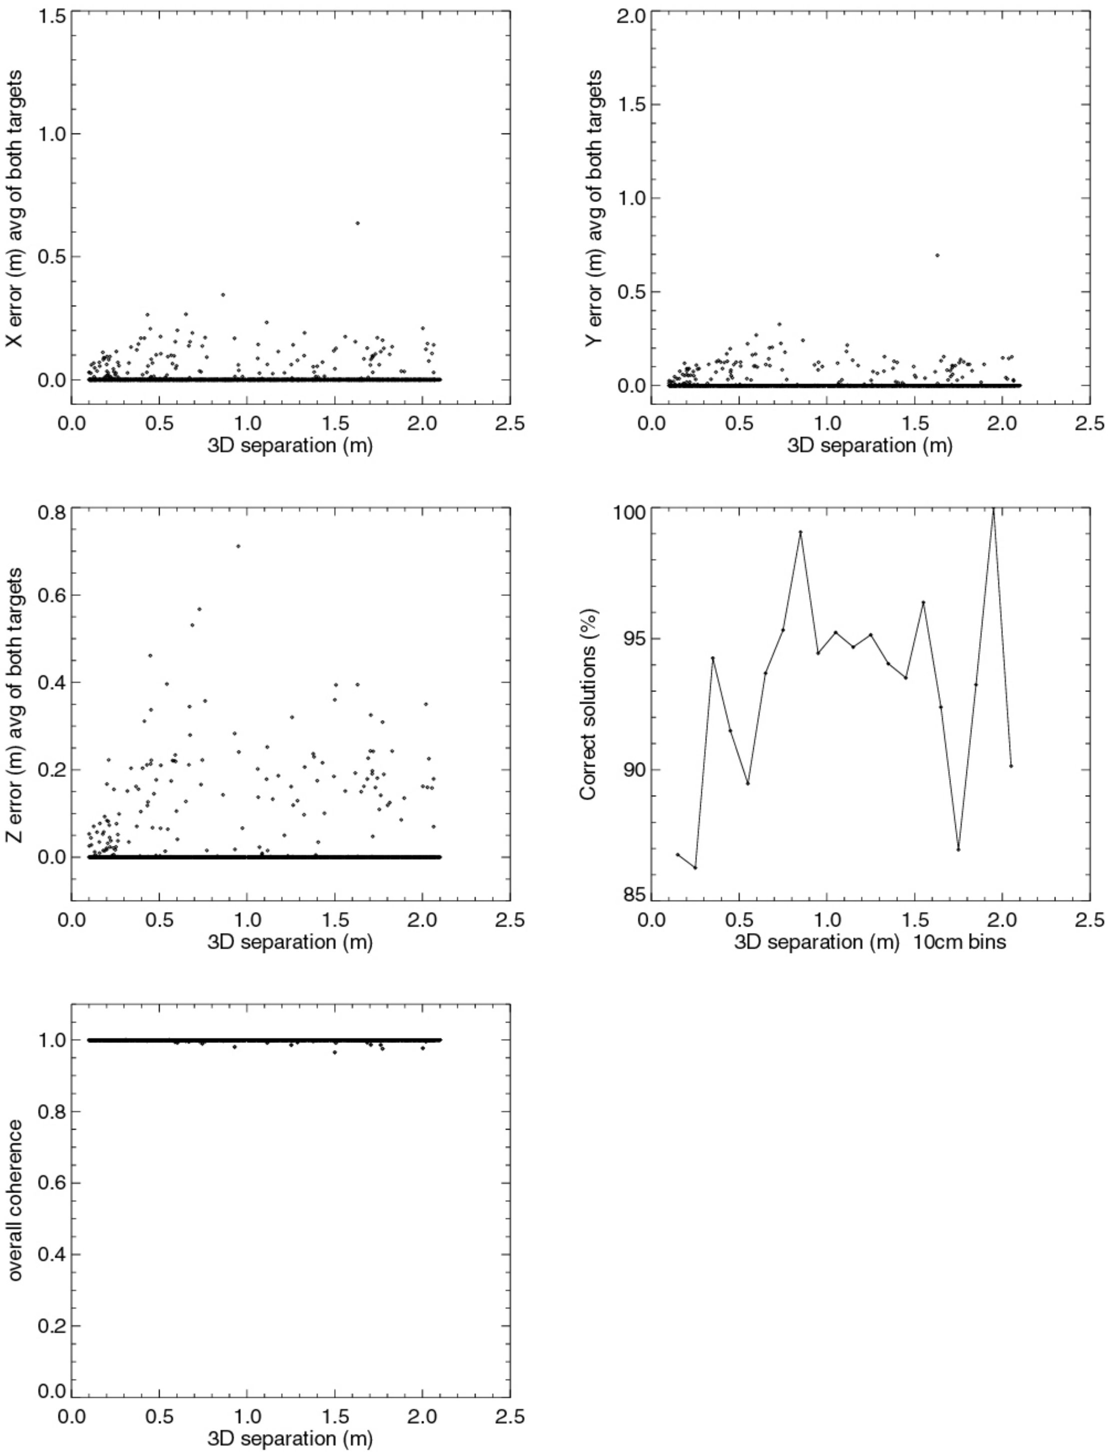

Figure 8. Double Happiness method on *electromagnetic* synthetic data, 2000 cases. All quantities are plotted versus separation in the *3D* plane.

5. Performance on Real Data

5.1 Data Description

Measurements of actual ordnance and clutter items were conducted by NRL at Blossom Point with a total field magnetometer, an EM61-MkII, an EM63 and a GEM-3. Four different types of ordnance were used (40mm, 60mm, 80mm, and 105mm), along with four pieces of actual clutter and two clutter clouds meant to simulate metal fragments. In addition, each of the targets was measured individually, with the ordnance positioned in three orientations (horizontal, vertical with nose up, and vertical with nose down). Finally, a background measurement was taken with no target present. Figures 9 and 10 show the library signatures obtained from these measurements.

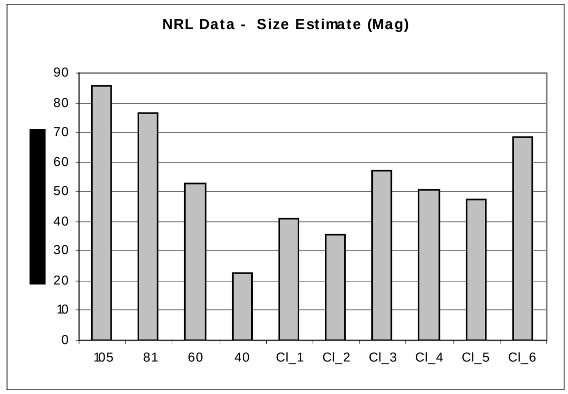

Figure 9. The magnetic sizes of the objects measured at Blossom Point.

In Figure 10, the plots show the longitudinal beta versus the average of the transverse betas for items measured singly. The "error bars" show the values of the transverse betas. Note that the ordnance has essentially equal transverse betas, implying a rod-like shape. The diagonal line represents equal longitudinal and transverse betas.

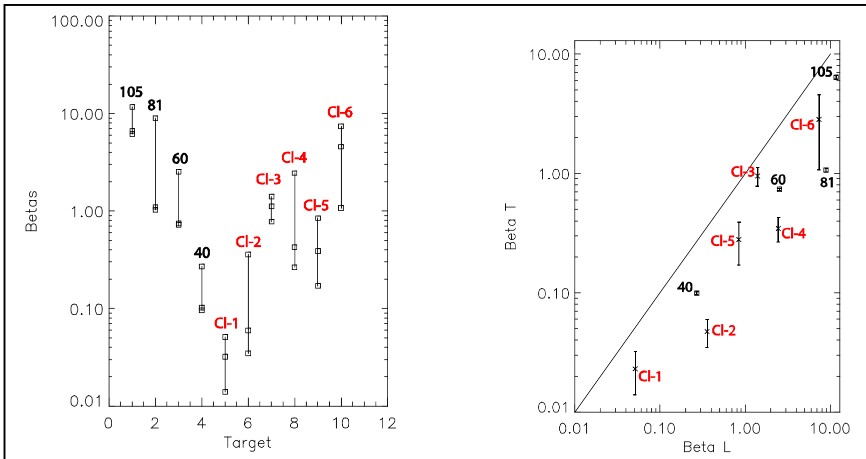

Figure 10. The beta values for the objects measured at Blossom

Approximately 70 different two-item combinations of these targets were measured, using various horizontal and vertical separations and target orientations. Figure 11 shows the distribution of 3D separations achieved.

Unfortunately, these measurements were obtained prior to the synthetic runs, so our growing knowledge of the problem area to investigate (i.e., small lateral separations) could not be taken into account when the data collection was planned. As a result, the database contains a range of target separations in depth but the lateral separations were quite small, with the vast majority having zero lateral separation. The largest lateral separation was 0.5m, and this occurred in only four cases. Indeed, in no case can the eye discern

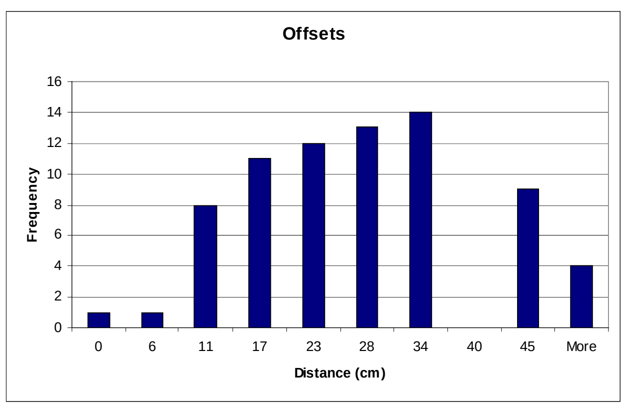

Figure 11. The distribution of 3D separations between object pairs measured at Blossom Point.

the presence of two targets from the measured signatures.

5.2 Results for the Iterative Residual Method

The iterative residual algorithm was used on 68 of the experimental setups from Blossom Point. The first step was to subtract the background from the measurements prior to fitting. Next, the overlapping targets were inverted using the iterative method and the fit parameters for the two resulting dipoles were then compared with the ground truth for the targets. As expected, the match was generally poor, with the EM results slightly better than the magnetic. There were two general outcomes. The majority of signatures were fit by one dipole, with a second, deep dipole acting as a perturbation but contributing little to the total signal. Effectively, then, in these cases the signal is being fit by a single dipole because the targets are too close to distinguish and, it is interesting to note, that the total fit coherence is often quite high for these false solutions. A minority of the fits contained two dipoles, nearly opposite in orientation, with minor differences in size and depth.

Figures 12 and 13 show the EM results for the four experiments with 0.5m horizontal separation of the objects. For these experiments, the same clutter item Cl-3 was paired with each of the four ordnance items placed horizontally. The plots show the same beta values as Figure 10, with the two specific targets used for each experiment circled. The inverted betas using the iterative method are plotted as diamonds and shaded yellow for ordnance and green for clutter. The ordnance betas are well retrieved for the 3 largest targets of the 4, and the clutter betas are also well retrieved in 3 of the 4 cases, although not for the same 3 cases.

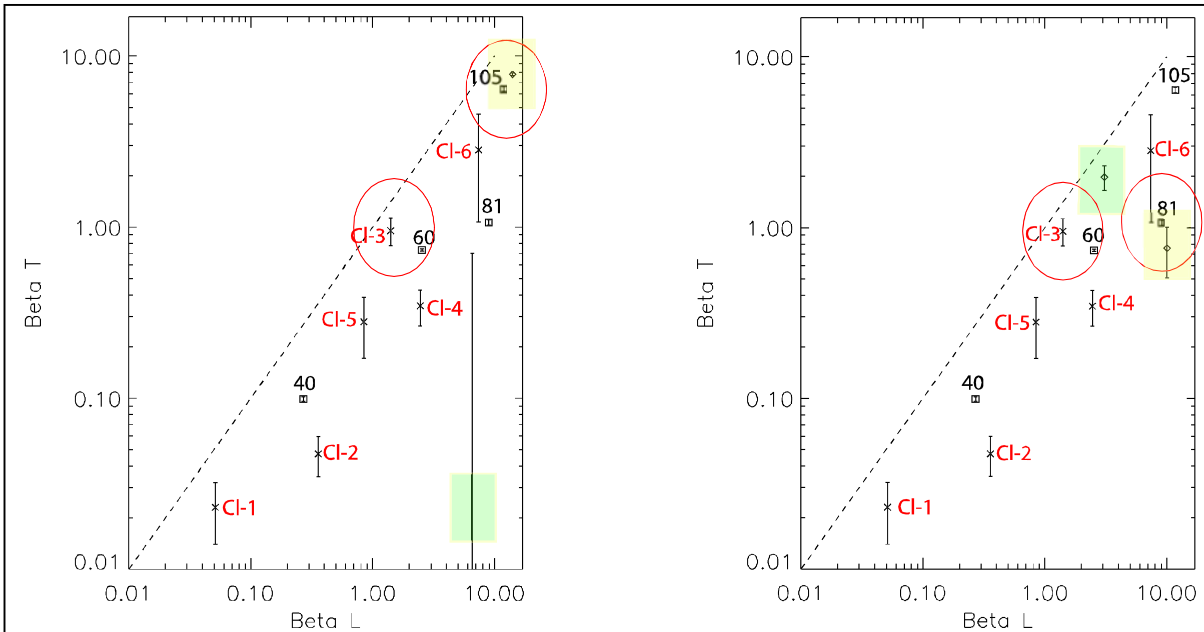

Figure 12. EM61MkII experiments 109 (left) and 110 (right), in which clutter item Cl-3 was paired with the 105mm ordnance (left) and the 81mm ordnance (right).

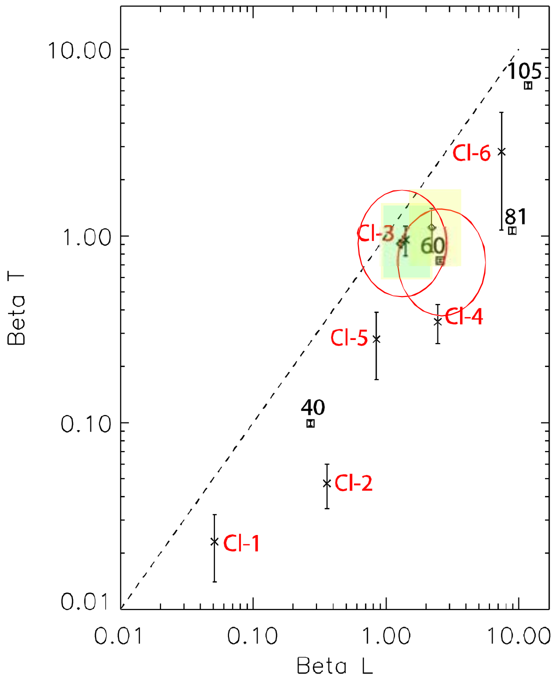

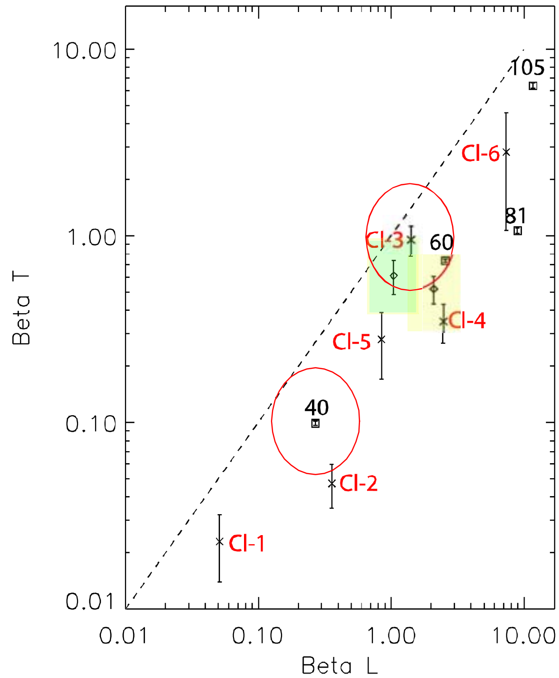

Figure 13. EM61MkII experiments 111 (left) and 112 (right), in which clutter item Cl-3 was paired with the 60mm ordnance (left) and the 40mm ordnance (right).

5.3 Results for Double Happiness Algorithm

The project was terminated before the final version of the double happiness algorithm could be run on the Blossom Point data.

6. Conclusions and Future Work

The original goal of this project was to develop advanced iterative techniques for inverting magnetic and electromagnetic data for situations in which the signatures from two targets overlap. The first method developed was the two-dipole iterative residual algorithm, a straightforward approach to the fit. A second technique was also developed - a simultaneous two-dipole fit algorithm we call “double happiness.”

On synthetic magnetometer data, the double happiness algorithm performed excellently, and better than the iterative residual method. On synthetic electromagnetic data, the double happiness algorithm still performed better than the iterative residual method, but differences in the assumed sensor (EM61MkI versus EM61MkII) make the direct comparison slightly ambiguous.

On overlapping signature data collected by NRL at Blossom Point with an EM61MkII, the iterative residual method performed reasonably well for the 4 cases of target pairs separated by 0.5m, but there was not enough data with real horizontal separations to adequately evaluate the algorithm.

Unfortunately, the project was terminated before we could test the final version of the double happiness algorithm on the data.

On the synthetic data it was seen that both algorithms had some difficulty separating two dipoles that are very close in horizontal separation, since the signature from the compact targets will appear like a single dipole. Fortunately, this situation is less important in practice than the situation in which a weaker target lies some distance from a stronger target, but the two signatures overlap and the weaker target is not detected. In such situations, traditional methods will locate the stronger target, but its excavation may miss the weaker one. In this case, the double happiness algorithm may perform more than adequately.

7. References

[1] Marquardt, "An Algorithm for Least-Squares Estimation of Nonlinear Parameters", J. Soc. Ind. Appl. Math., Vol 11, no. 2, pp. 431-441, June, 1963.